

Docket No. AT9-99-319

**METHOD AND SYSTEM FOR REMOTE JAVA AUDIO SERVER IN A
HETEROGENEOUS DISTRIBUTED ENVIRONMENT**

BACKGROUND OF THE INVENTION

5

1. Technical Field:

The present invention relates to an improved data processing system and, in particular, to a method and system for multicomputer data transferring, specifically, distributed data processing in a client-server environment.

2. Description of Related Art:

Given the increasing capabilities and decreasing costs of information processing and data networking technologies, the network is rapidly expanding its reach. The emerging infrastructure of computers interconnected by networks presents users with the ability to execute applications in a variety of configurations.

The X Window system is a network-based graphical windowing system for workstations. X Windows has been widely adopted as an industry standard supported by a consortium of industry leaders. Many Unix machines include X Windows as part of their operating system.

The X Window system is based upon the client-server model. A client application is an application program which performs a specific task. A display server application is a program which acts as the intermediary between client applications and the machine hardware. The display server application keeps track of all user and client application input/output. Ordinarily, the

Docket No. AT9-99-319

display server program is started automatically, or once started, executes for an entire user session.

One challenge presented to application developers by a network computing environment is the wide range of devices that are interconnected by networks. A typical network usually has many different kinds of attached devices with diverse hardware architectures, operating systems, and application support. Java runtime environments address this challenge by enabling the creation of platform-independent applications. A single Java application can execute unchanged on a wide variety of hardware platforms. Compared with applications written for a specific hardware and operating system platform, platform-independent programs written in Java can be easier and cheaper to develop and maintain. A Java application may run on a Java runtime platform that has been tailored for execution on a specific platform consisting of specific hardware and its platform-specific supporting operating system. By developing a Java runtime environment for various operating systems, the Java runtime environment acts as an intermediate layer between a Java application and a platform-specific operating system and provides platform-independence for the Java application.

Java runtime environments have been developed for several different computing platforms, such as Unix workstations. When a Java application is executing on a Unix host machine through X Windows, the graphics generated by the Java application may be distributed using the X Windows protocol to an X Windows server that is running on a client machine.

The Java runtime environment contains an audio playback engine that uses the native audio support of the

Docket No. AT9-99-319

supporting machine to play the audio data. However,
there is no support for audio in the X Windows protocol.
In most cases, when a Java application on a host machine
generates an audio datastream, the underlying
5 platform-specific operating system will employ the audio
support of the underlying machine so that the audio will
be audible at the underlying host machine and not the
client machine. Hence, if a user at a client machine
remotely executes a Java application on a host machine,
10 the graphics for the Java application will appear on the
user's client machine, but the audio will be played on
the remote host machine, and the user will not hear the
generated audio on the client machine. This situation is
possibly unexpected to the user and certainly
15 undesirable.

Therefore, it would be advantageous to provide a
method and system for remote audio processing in a
distributed heterogeneous environment, and in particular,
to provide distributed audio that has the advantages
20 provided by Java platform-independence.

2025 RELEASE UNDER E.O. 14176

A method and system for an audio server in a heterogeneous distributed environment is provided. A Java application executes on a host machine under X Windows or RAWT (Remote Abstract Window Toolkit), and the Java application generates audio data and graphic data. The graphic data is sent to a display server on a client machine specified by a display environment variable. Although neither X Windows nor RAWT have audio support, a Java audio driver on the host machine determines whether an audio environment variable or an audio command line parameter is specified on the host machine. In parallel to the graphic data, the audio data is then sent to a Java audio server on the client machine specified by the audio environment variable or the audio command line parameter, and played using the local audio support, in Java, on the client machine on which the user can hear the audio.

BRIEF DESCRIPTION OF THE DRAWINGS

5 The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself, however, as well as a preferred mode of use, further objectives and advantages thereof, will best be understood by reference to the following detailed
10 description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

Figure 1 depicts a pictorial representation of a distributed data processing system in which the present invention may be implemented;

15 **Figure 2A** is a block diagram of a data processing system which may be implemented as a server in accordance to the present invention;

Figure 2B is a block diagram of a data processing system in which the present invention may be implemented;

20 **Figure 3** is a block diagram illustrating the relationship of software components operating within a computer system that may implement the present invention

Figure 4 is a block diagram depicting a typical organization for a remote Java application supported by
25 an X Windows environment;

Figure 5 is a block diagram depicting a distributed audio server implemented in accordance with the present invention; and

30 **Figure 6** is a flowchart depicting a process for generating and playing audio data in a distributed data processing system in accordance with the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

5 With reference now to the figures, and in particular with reference to Figure 1, a pictorial representation of a distributed data processing system in which the present invention may be implemented is depicted.

10 Distributed data processing system 100 is a network of computers in which the present invention may be implemented. Distributed data processing system 100 contains a network 102, which is the medium used to provide communications links between various devices and computers connected together within distributed data
15 processing system 100. Network 102 may include permanent connections, such as wire or fiber optic cables, or temporary connections made through telephone connections.

In the depicted example, a server 104 is connected to network 102 along with storage unit 106. In addition,
20 clients 108, 110, and 112 also are connected to a network 102. These clients 108, 110, and 112 may be, for example, personal computers or network computers. For purposes of this application, a network computer is any computer, coupled to a network, which receives a program
25 or other application from another computer coupled to the network. In the depicted example, server 104 provides data, such as boot files, operating system images, and applications to clients 108-112. Clients 108, 110, and 112 are clients to server 104. Distributed data
30 processing system 100 may include additional servers, clients, and other devices not shown. In the depicted example, distributed data processing system 100 is the

Internet with network 102 representing a worldwide collection of networks and gateways that use the TCP/IP suite of protocols to communicate with one another. At the heart of the Internet is a backbone of high-speed data communication lines between major nodes or host computers, consisting of thousands of commercial, government, educational, and other computer systems, that route data and messages. Of course, distributed data processing system 100 also may be implemented as a number of different types of networks, such as, for example, an Intranet or a local area network.

Figure 1 is intended as an example, and not as an architectural limitation for the processes of the present invention.

With reference now to Figure 2A, a block diagram of a data processing system which may be implemented as a server, such as server 104 in Figure 1, is depicted in accordance to the present invention. Data processing system 200 may be a symmetric multiprocessor (SMP) system including a plurality of processors 202 and 204 connected to system bus 206. Alternatively, a single processor system may be employed. Also connected to system bus 206 is memory controller/cache 208, which provides an interface to local memory 209. I/O Bus Bridge 210 is connected to system bus 206 and provides an interface to I/O bus 212. Memory controller/cache 208 and I/O Bus Bridge 210 may be integrated as depicted.

Peripheral component interconnect (PCI) bus bridge 214 connected to I/O bus 212 provides an interface to PCI local bus 216. A modem 218 may be connected to PCI local bus 216. Typical PCI bus implementations will support four PCI expansion slots or add-in connectors.

Docket No. AT9-99-319

Communications links to network computers 108-112 in Figure 1 may be provided through modem 218 and network adapter 220 connected to PCI local bus 216 through add-in boards.

5 Additional PCI bus bridges 222 and 224 provide interfaces for additional PCI buses 226 and 228, from which additional modems or network adapters may be supported. In this manner, server 200 allows connections to multiple network computers. A memory mapped graphics
10 adapter 230 and hard disk 232 may also be connected to I/O bus 212 as depicted, either directly or indirectly.

Those of ordinary skill in the art will appreciate that the hardware depicted in Figure 2A may vary. For example, other peripheral devices, such as optical disk
15 drive and the like also may be used in addition or in place of the hardware depicted. The depicted example is not meant to imply architectural limitations with respect to the present invention.

The data processing system depicted in Figure 2A may
20 be, for example, an IBM RISC/System 6000 system, a product of International Business Machines Corporation in Armonk, New York, running the Advanced Interactive Executive (AIX) operating system.

With reference now to Figure 2B, a block diagram of
25 a data processing system in which the present invention may be implemented is illustrated. Data processing system 250 is an example of a client computer. Data processing system 250 employs a peripheral component interconnect (PCI) local bus architecture. Although the
30 depicted example employs a PCI bus, other bus architectures such as Micro Channel and ISA may be used. Processor 252 and main memory 254 are connected to PCI

Docket No. AT9-99-319

local bus 256 through PCI Bridge 258. PCI Bridge 258 also may include an integrated memory controller and cache memory for processor 252. Additional connections to PCI local bus 256 may be made through direct component interconnection or through add-in boards. In the depicted example, local area network (LAN) adapter 260, SCSI host bus adapter 262, and expansion bus interface 264 are connected to PCI local bus 256 by direct component connection. In contrast, audio adapter 266, graphics adapter 268, and audio/video adapter (A/V) 269 are connected to PCI local bus 266 by add-in boards inserted into expansion slots. Expansion bus interface 264 provides a connection for a keyboard and mouse adapter 270, modem 272, and additional memory 274. SCSI host bus adapter 262 provides a connection for hard disk drive 276, tape drive 278, and CD-ROM 280 in the depicted example. Typical PCI local bus implementations will support three or four PCI expansion slots or add-in connectors.

An operating system runs on processor 252 and is used to coordinate and provide control of various components within data processing system 250 in Figure 2B. The operating system may be a commercially available operating system such as JavaOS For Business™ or OS/2™, which are available from International Business Machines Corporation™. JavaOS is loaded from a server on a network to a network client and supports Java programs and applets. A couple of characteristics of JavaOS that are favorable for performing traces with stack unwinds, as described below, are that JavaOS does not support paging or virtual memory. An object oriented programming

Docket No. AT9-99-319

system such as Java may run in conjunction with the operating system and may provide calls to the operating system from Java programs or applications executing on data processing system 250. Instructions for the
5 operating system, the object-oriented operating system, and applications or programs are located on storage devices, such as hard disk drive 276 and may be loaded into main memory 254 for execution by processor 252. Hard disk drives are often absent and memory is
10 constrained when data processing system 250 is used as a network client.

Those of ordinary skill in the art will appreciate that the hardware in Figure 2B may vary depending on the implementation. For example, other peripheral devices,
15 such as optical disk drives and the like may be used in addition to or in place of the hardware depicted in Figure 2B. The depicted example is not meant to imply architectural limitations with respect to the present invention. For example, the processes of the present
20 invention may be applied to a multiprocessor data processing system.

The present invention provides a process and system for an audio server in a heterogeneous distributed environment. Although the present invention may operate
25 on a variety of computer platforms and operating systems, it may also operate within a Java runtime environment. Hence, the present invention may operate in conjunction with a Java virtual machine (JVM) yet within the boundaries of a JVM as defined by Java standard
30 specifications. In order to provide a context for the present invention, portions of the operation of a JVM according to Java specifications are herein described.

With reference now to Figure 3, a block diagram illustrates the relationship of software components operating within a computer system that may implement the present invention. Java-based system 300 contains platform specific operating system 302 that provides hardware and system support to software executing on a specific hardware platform. JVM 304 is one software application that may execute in conjunction with the operating system. JVM 304 provides a Java runtime environment with the ability to execute Java application or applet 306, which is a program, servlet, or software component written in the Java programming language. The computer system in which JVM 304 operates may be similar to data processing system 200 or computer 100 described above. However, JVM 304 may be implemented in dedicated hardware on a so-called Java chip, Java-on-silicon, or Java processor with an embedded picoJava core.

At the center of a Java runtime environment is the JVM, which supports all aspects of Java's environment, including its architecture, security features, mobility across networks, and platform independence.

The JVM is a virtual computer, i.e. a computer that is specified abstractly. The specification defines certain features that every JVM must implement, with some range of design choices that may depend upon the platform on which the JVM is designed to execute. For example, all JVMs must execute Java bytecodes and may use a range of techniques to execute the instructions represented by the bytecodes. A JVM may be implemented completely in software or somewhat in hardware. This flexibility allows different JVMs to be designed for mainframe computers and PDAs.

The JVM is the name of a virtual computer component that actually executes Java programs. Java programs are not run directly by the central processor but instead by the JVM, which is itself a piece of software running on the processor. The JVM allows Java programs to be executed on a different platform as opposed to only the one platform for which the code was compiled. Java programs are compiled for the JVM. In this manner, Java is able to support applications for many types of data processing systems, which may contain a variety of central processing units and operating systems architectures. To enable a Java application to execute on different types of data processing systems, a compiler typically generates an architecture-neutral file format - the compiled code is executable on many processors, given the presence of the Java runtime system. The Java compiler generates bytecode instructions that are nonspecific to a particular computer architecture. A bytecode is a machine independent code generated by the Java compiler and executed by a Java interpreter. A Java interpreter is part of the JVM that alternately decodes and interprets a bytecode or bytecodes. These bytecode instructions are designed to be easy to interpret on any computer and easily translated on the fly into native machine code. Bytecodes may be translated into native code by a just-in-time compiler or JIT.

A JVM must load class files and execute the bytecodes within them. The JVM contains a class loader, which loads class files from an application and the class files from the Java application programming interfaces (APIs) which are needed by the application. The execution engine that executes the bytecodes may vary across platforms and implementations.

One type of software-based execution engine is a just-in-time compiler. With this type of execution, the bytecodes of a method are compiled to native machine code upon successful fulfillment of some type of criteria for jitting a method. The native machine code for the method is then cached and reused upon the next invocation of the method. The execution engine may also be implemented in hardware and embedded on a chip so that the Java bytecodes are executed natively. JVMs usually interpret bytecodes, but JVMs may also use other techniques, such as just-in-time compiling, to execute bytecodes.

When an application is executed on a JVM that is implemented in software on a platform-specific operating system, a Java application may interact with the host operating system by invoking native methods. A Java method is written in the Java language, compiled to bytecodes, and stored in class files. A native method is written in some other language and compiled to the native machine code of a particular processor. Native methods are stored in a dynamically linked library whose exact form is platform specific.

The present invention provides a process and system for an audio server in a heterogeneous distributed environment. The audio server provides an important link for a Java application running in a heterogeneous distributed environment in which the host system may be a larger multi-user system and the user system may be a desktop computer. A complete range of Java audio functionality is thus provided between the systems connected through the network. Although the following figures and examples describe a distributed audio server that employs the capabilities of a Java runtime environment operated in conjunction with an X Windows

Docket No. AT9-99-319

distributed environment or RAWT (Remote Abstract Window Toolkit) running on any hardware platform, the present invention may also be implemented using other runtime environments or operating systems. In other words, Java provides platform-independence for the audio server, and as examples, RAWT or X Windows provides a platform-independent or platform-specific environment, respectively, for the distributed functionality of a display server. However, other platform-independent and platform-specific software could provide similar capabilities so that one may develop the distributed audio server of the present invention in other environments. The supporting environments are intended as examples and not as architectural limitations for the processes of the present invention.

On a client machine operating under an x windows server, the X Windows software uses a DISPLAY environment variable in order to connect the user's client machine with the X Windows server. The DISPLAY environment variable is related to the Internet protocol (IP) address of the client machine. Each client machine has a different IP address, and the IP address has the form "xxx.xxx.xxx.xxx", where "xxx" represents an 8-bit value. The IP address uniquely identifies the client machine for a variety of software and network purposes.

The DISPLAY environment variable provides a label usable within application code for addressing a display to which an application may connect. In this manner, the application code may reference the DISPLAY environment variable without hard-coding a pointer to a specific display within the application code.

Environment variables generally reside in configuration files on a client machine. In the most

Docket No. AT9-99-319

common case, a user may set an environment variable within the configuration file in order to pass configuration values to an application program. For example, in the DOS operating system, a user may set environment variables within the "autoexec.bat" file, or the .profile or .dtprofile on AIX, and environment variables within the file are then available in the runtime environment for various applications. The operating system may read the values in the configuration file in order to initialize the environment variables upon startup, or an application program may access the configuration file to dynamically read a value for an environment variable.

In the Java runtime environments, values may be stored in property files that are associated with Java applications. In another example, the Microsoft Windows operating system provides registry files that may be used to set environment variables for various applications. In some Unix implementations, a DISPLAY environment variable may be stored in a user's .login file or .profile file.

A value for a DISPLAY environment variable is generally specified using an assignment statement such as "DISPLAY=xxx.xxx.xxx.xxx". In this manner, the environment variable contains the IP address to which graphic data should be directed for display output, and the value may be changed to an appropriate IP address without changing the source code within an application program. In an x windows system, the DISPLAY environment variable would specify the IP address of the machine that is running an x windows server. A command line parameter may be specified by a user, however, in order to dynamically point the output of an application to a

Docket No. AT9-99-319

specific IP address. The command line parameter would take precedence over the environment variable specified in a file.

With reference now to Figure 4, a block diagram depicts a typical organization for a remote Java application supported by an x windows environment. Host machine 400 and client machine 402 are connected by network link 404. User Java application 406 executes on host machine 400 and sends graphic data via network link 404 to X Windows server 408 executing on client machine 402. X Windows 408 then directs graphic data to display 410. DISPLAY environment variable 412 is stored within configuration file 414. The DISPLAY environment variable indicates the IP address of client machine 402 to which host machine 400 should direct graphic data generated by applications executing on host machine 400.

With reference now to Figure 5, a block diagram depicts a distributed audio server implemented in accordance with the present invention. Host machine 500 and client machine 502 are connected by network link 504. User Java application 506 executes on host machine 500 and sends graphic data via network link 504 to X Windows server 508 executing on client machine 502. X Windows 508 then directs graphic data to display 510. DISPLAY environment variable 512 is stored within configuration file 514. The DISPLAY environment variable indicates the IP address of client machine 502 to which host machine 500 should direct graphic data generated by applications executing on host machine 500.

Figure 5 is similar to Figure 4 in that a platform-specific environment, such as X Windows,

Docket No. AT9-99-319

supports a platform-independent environment, such as Java. However, Figure 5 shows the manner in which an environment variable may be used to direct the audio generated by a user application to a client machine in use by the user. In the present invention, a special environment variable is used to direct audio data across the network to the Java audio server. Configuration file 514 contains AUDIO_DISPLAY environment variable 516 that specifies an IP address and port number to which host machine 500 should direct audio data generated by user Java application 506. Alternatively, a command line parameter may be specified by a user, however, in order to dynamically point the output of an application to a specific IP address. The command line parameter would take precedence over the environment variable specified in a file. Once the reference to the client machine has been retrieved, however, the present invention operates in the same manner no matter what the source of the reference.

The AUDIO_DISPLAY environment variable may be specified in a configuration file or properties file using an assignment statement such as "AUDIO_DISPLAY=xxx.xxx.xxx.xxx:P". Client machine 502 receives the audio data via network link 505 and passes the audio data to the application connected to the specified port number. In this case, Java audio server 518 processes the audio data and uses the underlying platform to generate sound through speaker 520. Network link 504 and network link 505 may be different logical network connections on the same physical network. Alternatively, client machine 502 may receive the audio data via different physical communication links.

In the preferred embodiment, the supporting code on the host machine just below the audio mixer writes audio buffers to an audio driver or multimedia subsystem on the host machine. The audio driver is preferably implemented
5 in the Java runtime at the application level, thereby isolating it from affecting other applications in the system. Alternatively, the audio driver is implemented as a device driver as part of the operating system.

The audio driver on the host machine checks for the
10 existence of a specified environment variable (or command line parameter), and if a special environment variable is defined, the audio driver writes the audio data with the appropriate protocol header information to the IP address specified in the AUDIO_DISPLAY environment variable.

15 The audio data may be broken into small data packets and sent to the client machine using the User Datagram Protocol (UDP). UDP is a connectionless protocol within TCP/IP that corresponds to the transport layer in the ISO network model. UDP converts data messages generated by
20 an application into packets to be sent via the IP protocol but does not verify that messages have been correctly delivered. Alternatively, the Real Time Streaming (RTP) protocol may be used for the audio data packets as this protocol builds upon UDP.

25 The audio data packets may arrive at the audio server on the client machine in a timely manner, but because of the methodology of UDP, these packets may sometimes be out of order or a packet may be missing depending upon a network load. The audio server on the
30 client machine will queue the packets and reorder the packets in order to play the audio content using the native audio support available on the client machine. Audio format processing and mixing may occur on the host

machine in order to minimize the network bandwidth.

The Java-based audio server may be written in pure Java and may utilize the audio capabilities of the client machine. In this manner, the audio data may be played
5 for the user at the client machine as if the audio data has been generated locally. The client machine need not be the same type of machine as the host machine given the platform-independence provided by Java.

When distributing audio data, it is important to
10 provide some synchronization support so that other media may be played in a corresponding way, such as video data. In this implementation, when audio is sent from the host machine to the audio server on the client machine, it should be considered real-time data. Currently, in the
15 Java Media Frameworks, video is synchronized with audio at the application level by writing audio packets to the audio device as near real-time as possible. In the present invention, the packets may be sent over the Internet and then played on the audio server at the
20 user's machine. There will naturally be some small delay, but because the system may use UDP, the performance provided by UDP should be sufficient in most cases for adequate audio quality. The audio server may periodically send back some type of information to the
25 generating application that provides an estimate of the delay time between the time point when the audio packet was sent and the time point at which the content within the audio packet was played. This estimate will provide the generating application with a better capability for
30 synchronizing the generated audio and video that is sent to the client machine.

With reference now to Figure 6, a flowchart depicts a process for generating and playing audio data in a

distributed data processing system in accordance with the present invention. The process begins with the generation of audio and graphic data by a platform-independent user application, such as a Java application (step 602). A platform-specific graphics environment receives the graphic data and sends the graphics to a display server on the user's client machine in accordance with a display environment variable (step 604). For example, an x windows environment on the host machine accepts the graphics and send the graphics to an x windows display server on the client machine in accordance with a DISPLAY environment variable, which may include displaying the graphic data on the host machine if the environment variable so dictates.

An audio driver on the host machine receives the audio data (step 606) and determines whether an environment variable (or command line parameter) specifies a specific machine for playing the audio data (step 608). If not, then the audio driver plays the audio data on its local machine (step 610), and the process completes. If an audio environment variable is specified, then the audio driver sends the audio data to the location or address specified by the audio environment variable (step 612). A platform-independent audio server receives the audio data (step 614) and plays the audio data on the client machine (step 616). The process is then completed with respect to the distributed, platform-independent audio server.

The advantages provided by the present invention should be apparent in view of the detailed description of the invention provided above. The present invention provides a distributed audio server that may be executed

in a platform-independent manner on any client machine that supports the distributed execution or remote execution of an application. This allows audio data that is generated by a remote platform-independent application to be distributed to a user's client machine in parallel with the graphic data generated by the same application so that the audio content will be audible by the user.

It is important to note that while the present invention has been described in the context of a fully functioning data processing system, those of ordinary skill in the art will appreciate that the processes of the present invention are capable of being distributed in the form of a computer readable medium of instructions and a variety of forms and that the present invention applies equally regardless of the particular type of signal bearing media actually used to carry out the distribution. Examples of computer readable media include recordable-type media such a floppy disc, a hard disk drive, a RAM, and CD-ROMs and transmission-type media such as digital and analog communications links.

The description of the present invention has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to the invention in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art. The embodiment was chosen and described in order to best explain the principles of the invention, the practical application, and to enable others of ordinary skill in the art to understand the invention for various embodiments with various modifications as are suited to the particular use contemplated.